

Robustness and efficiency of covariate adjusted linear instrumental variable estimators

Stijn Vansteelandt

*Department of Applied Mathematics, Computer Sciences and Statistics
Ghent University, Belgium*

and Vanessa Didelez
*School of Mathematics
University of Bristol, U.K.*

Two-stage least squares (TSLS) estimators and variants thereof are widely used to infer the effect of an exposure on an outcome using instrumental variables (IVs). They belong to a wider class of two-stage IV estimators, which are based on fitting a conditional mean model for the exposure, and then using the fitted exposure values along with the covariates as predictors in a linear model for the outcome. We show that standard TSLS estimators enjoy greater robustness to model misspecification than more general two-stage estimators. However, by potentially using a wrong exposure model, e.g. when the exposure is binary, they tend to be inefficient. In view of this, we study double-robust G-estimators instead. These use working models for the exposure, IV and outcome but only require correct specification of either the IV model or the outcome model to guarantee consistent estimation of the exposure effect. As the finite sample performance of the locally efficient G-estimator can be poor, we further develop G-estimation procedures with improved efficiency and robustness properties under misspecification of some or all working models. Simulation studies and a data analysis demonstrate drastic improvements, with

remarkably good performance even when one or more working models are misspecified.

Key-words: bias; confounding; instrumental variable; model misspecification; semi-parametric efficiency.

1 Introduction

An enormous body of research has developed in the econometrics and biostatistics literatures on how to assess the causal effect of an exposure X on an outcome Y in the presence of confounding by unobserved variables U , when a *vector of instrumental variables* Z (IVs) is available (see e.g. Bowden and Turkington, 1985; Robins, 1994; Angrist et al., 1996; Greenland, 2000; Wooldridge, 2002; Hernán and Robins, 2006; Didelez and Sheehan, 2007). It is therefore not surprising that a variety of competing approaches have been put forward. A simple and popular method is two-stage least squares (TSLS) estimation where, in the first stage, the exposure is predicted based on an ordinary least squares regression of the exposure on the IVs and covariates; in the second stage, the outcome is regressed on the predicted exposure and covariates via ordinary least squares regression, and the exposure coefficient is taken as the final IV estimator of the desired causal effect. The simplicity of this approach has encouraged the development of other two-stage estimators, which are obtained along the same lines, but employ possibly non-linear regressions in the first or second stage (see e.g. Mullahy, 1997; or review in Didelez et al., 2010). Variations on two-stage estimators, that we do not consider in much detail here, include limited information maximum likelihood (LIML; Bowden and Turkington, 1985, chapter 4; or see Anderson 2004 for an historic account), Bayesian (Kleibergen and Zivot, 2003) or control function approaches (Wooldridge, 2002, chapter 6). In view of their increasing popularity, we will contrast the two-stage approaches with G-estimators under structural mean models. These do not rely on separate fitting of a first stage model

(Hernán and Robins, 2006) and have connections to Generalized Method of Moments (GMM) estimation (Clarke and Windmeier, 2010).

It is natural to ask how the various IV methods compare with regard to their efficiency as well as robustness under various types of model misspecification. In particular, two-stage methods appear to rely on a correct exposure model but may in a variety of situations be consistent even if this is misspecified, or not be efficient even if the exposure model is correct. Moreover, in the presence of covariates, one can investigate the role of models for the relation between covariates on the one hand, and exposure, outcome or instruments on the other hand. Efficiency is relevant in this context as methods that are robust towards misspecification of some model assumptions will likely not be as efficient as methods exploiting a correctly specified model. Moreover, one can ask whether including covariates (when there is the choice) typically leads to efficiency gains (as noticed in an IV setting by Fisher-Lapp and Goetghebeur, 1999) and whether there is a trade-off with robustness.

In this article we investigate these questions formally in the context of linear IV models; these are introduced in Section 2. We then focus on two-stage IV estimators where we consider arbitrary (possibly nonlinear) conditional mean models for the exposure combined with arbitrary linear conditional mean models for the outcome (Section 3). We study their efficiency, and their bias under misspecification of the exposure and/or outcome model. In Section 4 we find a subclass of two-stage IV estimators to enjoy robustness against misspecification of the exposure model. We moreover derive the locally efficient IV-estimator which does not rely on correct specification of an exposure model, and find it to equal a specific two-stage estimator in some cases, but not in general. As addressed in Section 5, a further subclass of IV estimators is double-robust: consistent if either a model for the main effect of covariates on the outcome or a model for the distribution of the IV, given covariates, is correctly specified, but not necessarily both

(Okui et al., 2012). Interestingly, the TSLS estimator enjoys double robustness, but only w.r.t. a linear model for the conditional mean of the IV given covariates (Robins, 2000). Moreover, it is sometimes inefficient relative to other double-robust estimators, for instance when the exposure obeys a nonlinear model or when the exposure effect depends on covariates.

In our simulation studies the locally efficient double-robust IV estimator (Robins, 1994) outperforms the TSLS estimator when based on correctly specified exposure and outcome models, but performs much worse otherwise. In view of this, we develop two adaptive estimation procedures in Section 6. The first makes use of empirical efficiency maximisation (Rubin and van der Laan, 2008) which is designed to maximise precision even under misspecification of the exposure and outcome model, and results in drastic efficiency gains when the model for the distribution of the IV, given covariates, is correctly specified. The second makes use of bias-reduced double-robust estimation (Vermeulen and Vansteelandt, 2015), which is designed to minimise bias even under additional misspecification of the IV distribution. Simulation studies confirm the bias-reduction and moreover demonstrate favourable performance regarding efficiency.

Proofs of the various propositions and corollaries in the article can be found in Appendix A of the Supplemental Materials.

2 Linear instrumental variable models

Let Z be a vector of IVs for the effect of a scalar exposure X on a scalar outcome Y , conditional on a vector of observed covariates C . In analogy to Didelez and Sheehan (2007) but extending the definition to account for covariates (see also Pearl, 2009, p.248), we formalise this by the assumptions that Z is (a) associated with X conditional on C , (b) independent of Y , conditional on X, U and C , and (c) independent of U , conditional on

C ; here, U is a (set of) latent variable(s) such that (U, C) would be sufficient to control for confounding of the effect of X on Y were U observable. This formalisation of an IV is close to, but allows for greater flexibility than that in the econometric literature on IVs (Wooldridge, 2002), where assumptions are usually in terms of no correlation instead of independence. In the causal inference literature, conditions (b) and (c) are often alternatively formalised in the assumption (b') that $Y_x \perp\!\!\!\perp Z|C$ (Robins, 1994), with Y_x denoting the counterfactual outcome that would be observed when setting X to x . The latter formulation avoids explicit reference to any specific unobserved confounders.

We start by briefly addressing the relationship of different formulations of linear IV models. Consider first the following model for the conditional mean of the outcome:

$$E(Y|X, Z, U, C) = \omega(C, U) + m(C; \psi^*)X. \quad (1)$$

Here, $\omega(C, U)$ is an unknown (i.e., unspecified) function of measured and unmeasured covariates. The term $m(C; \psi)$ is a known function of observed covariates, smooth in ψ , and ψ^* is an unknown finite-dimensional parameter, e.g. $m(C; \psi) = \psi$ or $m(C; \psi) = \psi^T C$, where with a slight abuse of notation, the vector C includes 1 to allow for a main effect. When $m(C; \psi)$ is parameterised such that $m(C; \psi) = 0$ when $\psi = 0$, as in the previous examples, we have that $m(C; \psi^*)$ captures the exposure effect of interest, i.e. we regard ψ^* as the target ‘causal’ parameter. In particular,

$$\begin{aligned} m(C; \psi^*)X &= E(Y|X, Z, U, C) - E(Y|X = 0, Z, U, C) \\ &= E(Y|X, Z, U, C) - E(Y_0|X = 0, Z, U, C) \\ &= E(Y|X, Z, U, C) - E(Y_0|X, Z, U, C) \\ &= E(Y|X, Z, C) - E(Y_0|X, Z, C), \end{aligned} \quad (2)$$

which encodes the additive effect on the outcome of setting the exposure to zero in a subgroup of individuals with exposure X , IV Z and covariates C . In the above derivation,

the second equality follows by the consistency assumption that $Y = Y_0$ in subjects with $X = 0$, the third from the fact that U and C are sufficient to control for confounding of the effect of X on Y , and the fourth by the fact that the left-hand side does not involve U . Note that many IV methods assume a (linear) structural equation for the outcome Y which is more restrictive but implies the above (1) for $\omega(C, U)$ equal $\beta^{*T}C + U$.

The model defined by restriction (2), i.e.

$$E(Y|X, Z, C) = E(Y_0|X, Z, C) + m(C; \psi^*)X, \quad (3)$$

is called a linear or additive structural mean model (Robins, 1994). Together with the IV assumptions (a) and (b') it can be regarded as the substantive model of interest, as it merely parameterizes the exposure effect of interest. That the exposure effect does not involve Z (or equivalently, that Z does not appear on the right-hand side of (1), but is included on the left hand side) is known as ‘no effect modification’ by Z (Hernan and Robins, 2006; Clarke and Windmeier, 2010). It is this assumption which ultimately allows inference exploiting the IV Z as we will see below. While it can be motivated by the additivity in (1), it is often made in its own right avoiding explicit reference to U and hence allowing greater generality. In other words, models (1) and (3) differ in their assumptions on unobservables, but they essentially impose the same restrictions on the observed data law under the IV assumptions (see the Supplementary Materials for details). We therefore use the same notation, \mathcal{M} , throughout to denote both IV models.

3 Two-stage estimation

Model \mathcal{M} cannot be fitted directly as Y_0 (resp. U) is unobserved. Two-stage approaches exploiting the IV Z use the following restriction implied by \mathcal{M} :

$$E(Y|Z, C) = \omega(C) + m(C; \psi^*)E(X|Z, C), \quad (4)$$

for $\omega(C) \equiv E(Y_0|C)$ (or $\omega(C) \equiv E\{\omega(C, U)|C\}$ if we start with (1)). When C is high-dimensional (e.g. continuous or discrete with several components), the above cannot be fitted non-parametrically and additional modelling assumptions are needed to obtain estimators of ψ^* with adequate performance in moderate sample sizes. Equation (4) suggests postulating two additional models, one for $E(X|Z, C)$ and one for $\omega(C)$, and thereby lays the basis of two-stage estimation procedures.

In the first stage, a parametric model \mathcal{A}_x is postulated for the exposure, i.e.

$$E(X|Z, C) = m_x(Z, C; \alpha^*), \quad (5)$$

where $m_x(Z, C; \alpha)$ is a known function of instruments and covariates, smooth in α and α^* is an unknown finite-dimensional parameter. An obvious choice would be a linear or logistic regression model (e.g., $m_x(Z, C; \alpha) = \text{expit}(\alpha_z^T Z + \alpha_c^T C)$). The second stage model supplements that structural model \mathcal{M} with a parametric model \mathcal{A}_y for the main effects of covariates on the outcome:

$$\omega(C) = m_y(C; \beta^*), \quad (6)$$

where $m_y(C; \beta)$ is a known function of covariates, smooth in β and β^* is an unknown finite-dimensional parameter.

In the remaining sections, we will highlight results and computations that are specific to the following common choices for the models. We denote the model with a constant causal effect as

$$\mathcal{M}^{const} : \quad m(C; \psi) = \psi. \quad (7)$$

This is to be contrasted with the more general case allowing effect modification by some of the covariates in C . Further, we will pay special attention to linear models for exposure and covariates

$$\mathcal{A}_x^{lin} : \quad m_x(Z, C; \alpha) = \alpha_z^T Z + \alpha_c^T C \quad (8)$$

$$\mathcal{A}_y^{lin} : \quad m_y(C; \beta) = \beta^T C, \quad (9)$$

where it will often be important that the vector of covariates C is indeed exactly the same in both (8) and (9). Note that model \mathcal{A}_x^{lin} is quite general. In particular, it allows for covariate-instrument interactions by letting Z equal Z^*C for some IV Z^* . More generally, the two models could use the covariates in a different way, e.g. with V a scalar component of C , we may have $\alpha_1 + \alpha_2 \log V$ in modification of (8) combined with $\beta_1 + \beta_2 V + \beta_3 V^2$ in modification of (9).

3.1 Consistency of general two-stage estimators

The model defined by all three types of restrictions, i.e. $\mathcal{M} \cap \mathcal{A}_x \cap \mathcal{A}_y$, implies a standard conditional mean model (Chamberlain, 1987) of the form

$$E(Y|Z, C) = m_y(C; \beta^*) + m(C; \psi^*)m_x(Z, C; \alpha^*). \quad (10)$$

For instance, the models $\mathcal{M}^{const}, \mathcal{A}_x^{lin}, \mathcal{A}_y^{lin}$ imply the linear regression model

$$E(Y|Z, C) = \beta^{*T} C + \psi^*(\alpha_z^{*T} Z + \alpha_c^{*T} C).$$

Likewise, the logistic exposure model combined with a linear outcome model implies

$$E(Y|Z, C) = \beta^{*T} C + \psi^* \text{expit}(\alpha_z^{*T} Z + \alpha_c^{*T} C).$$

A general two-stage procedure is thus obtained by fixing α^* at some estimate $\hat{\alpha}$ obtained from fitting model (5) and then fitting model (10) with α^* substituted by $\hat{\alpha}$ using standard regression techniques at each stage. Two-stage estimation thus arises very naturally in an IV context. However, as we will show, it is not generally consistent under misspecification of model \mathcal{A}_x or \mathcal{A}_y , and not necessarily efficient even when both models are correct; we would therefore not generally recommend this approach. We cover two-stage estimation

here because some special cases, in particular the popular TSLS, exhibit greater robustness and efficiency, which we will later compare with other methods that are robust by design.

Recall that all Consistent Asymptotically Normal (CAN) estimators for α^* in \mathcal{A}_x are asymptotically equivalent to the solution to an estimating equation of the form

$$0 = \sum_{i=1}^n e_x(Z_i, C_i) \{X_i - m_x(Z_i, C_i; \alpha)\} \quad (11)$$

for index functions $e_x(Z, C)$ of the dimension of α . For example, when \mathcal{A}_x^{lin} is fitted using ordinary least squares estimation we have $e_x(Z, C) = (Z^T, C^T)^T$ in (11). Similarly, for a given α^* , all CAN estimators for β^*, ψ^* in model (10) are asymptotically equivalent to the solution to an estimating equation of the form

$$0 = \sum_{i=1}^n e_y(Z_i, C_i) \{Y_i - m_y(C_i; \beta) - m(C_i; \psi)m_x(Z_i, C_i; \alpha^*)\} \quad (12)$$

for index functions $e_y(Z, C)$ of the dimension of $(\beta^T, \psi^T)^T$. Under \mathcal{A}_y^{lin} and \mathcal{M}^{const} , using ordinary least squares estimation amounts to choosing $e_y(Z, C) = (C^T, \alpha_z^{*T}Z + \alpha_c^{*T}C)^T$ in (12). The estimators $\hat{\beta}$ for β^* and $\hat{\psi}$ for ψ^* resulting from substituting α^* by $\hat{\alpha}$ in (12) are also sometimes called ‘plug-in’ estimators; they are still CAN, as the following proposition asserts.

Proposition 1 Two-stage estimator is CAN

The two-stage IV estimator of the causal parameter ψ^ obtained by fixing α^* at the consistent estimator $\hat{\alpha}$ obtained by solving (11) for some conformable index function $e_x(Z, C)$, and next fitting model (10) by solving (12) for some conformable index function $e_y(Z, C)$ (with α^* substituted by $\hat{\alpha}$), is consistent and asymptotically normal under $\mathcal{M} \cap \mathcal{A}_x \cap \mathcal{A}_y$.*

Proof: this follows from the general theory of M-estimation under standard conditional mean models (Stefanski and Boos, 2002). \square

Although two-stage estimators are also sometimes used in nonlinear models for the outcome (e.g. logistic regression models and proportional hazard models), Proposition 1 relies on the identity (4) and is difficult to justify outside the realm of such additive causal models as (1) or (3) (see also Didelez, Meng and Sheehan, 2010; Vansteelandt et al., 2011); an exception occurs in the context of additive hazard models as discussed in Tchetgen Tchetgen et al. (2015).

3.2 Two-stage least-squares estimation (TSLS)

Among IV methods, TSLS takes a prominent place and provides an apparent two-stage estimator (Wooldridge, 2002). The principle of TSLS is that all ‘endogenous’ exposures (those that are confounded, i.e. dependent on U) are replaced by their linear projections on all ‘exogenous’ variables (these are the IVs, covariates, and possible other unconfounded exposures in the outcome model). As the name suggests, TSLS is equivalent to explicit two-stage estimation because the linear projections are equivalent to fitting a linear first stage model with ordinary least squares, and these can then be plugged into the second stage model, again fitted by least-squares. For the equivalence it is, however, important to use the *implied* first stage model, i.e. a linear model for X (and other endogenous variables) given *all* exogenous variables as determined by the choice of IVs and \mathcal{A}_y . In the following we refer to this procedure as *Standard TSLS*, to distinguish it from *Plug-In TSLS*, which also uses least squares in the first and second stage, but under possibly more general linear models $\mathcal{A}_x, \mathcal{A}_y$, e.g. with different transformations of the covariates.

We illustrate Standard TSLS with an example. Consider the case where $C = (1, V)^T$ with V a scalar, $m(C; \psi)X = \psi_1 X + \psi_2 XV$ and $m_y(C; \beta) = \beta_0 + \beta_1 V + \beta_2 V^2$. There are two ‘endogenous’ variables, X and XV , as these both depend on U . For identification it is necessary that there are at least as many instruments as endogenous variables; hence,

two instruments are needed, which could be Z and ZV . The linear projections would be of X and XV each on all of Z , ZV , V and V^2 . It follows that the implied first stage models are $E(X|Z, C) = \alpha_0 + \alpha_1 Z + \alpha_2 ZV + \alpha_3 V + \alpha_4 V^2$ and $E(XV|Z, C) = \alpha'_0 + \alpha'_1 Z + \alpha'_2 ZV + \alpha'_3 V + \alpha'_4 V^2$ where it is assumed that the coefficients of the instruments in the projections are non-zero (more precisely, that the matrix with first row α_1, α_2 and second row α'_1, α'_2 has full rank); the latter is a more specific version of assumption (a).

3.3 Efficiency of two-stage IV estimation

Even when model $\mathcal{M} \cap \mathcal{A}_x \cap \mathcal{A}_y$ is correctly specified, two-stage IV estimators are not generally efficient under this model because they are based on *separately* fitting the exposure model and the outcome model. Since the parameter α^* indexing the exposure model also appears in the outcome model (10), simultaneous fitting of the exposure and outcome model may sometimes result in more efficient estimators of ψ^* under model $\mathcal{M} \cap \mathcal{A}_x \cap \mathcal{A}_y$ (this is outlined in Appendix A of the Supplementary Materials as part of the proof of Proposition 2). However, Proposition 2 shows that efficiency is achieved for the Standard TSLS estimators, despite their apparent two-stage nature, under model $\mathcal{M}^{\text{const}} \cap \mathcal{A}_x^{\text{lin}} \cap \mathcal{A}_y^{\text{lin}}$, provided that the conditional variance-covariance matrix $\text{Cov}((X, Y)^T | Z, C)$ is constant in Z, C .

Proposition 2 Efficiency of Standard TSLS estimators

When the conditional variance-covariance matrix $\text{Cov}((X, Y)^T | Z, C)$ is constant in Z, C , then the Standard TSLS estimator of ψ^ is semi-parametric (locally) efficient in model $\mathcal{M}^{\text{const}} \cap \mathcal{A}_x^{\text{lin}} \cap \mathcal{A}_y^{\text{lin}}$.*

Proposition 2 does not immediately extend to more general two-stage IV estimators. In particular, two-stage estimators (including TSLS estimators) may be inefficient when

the exposure effect depends on covariates (i.e. when $m(C; \psi^*) = \psi^{*T}C$), even when the exposure and outcome model are fitted using ordinary least squares regressions. For TSLS estimators, this can be intuitively seen because for instance when $m(C; \psi^*) = \psi_0^* + \psi_1^*V$ for $C = (1, V)^T$, TSLS is based on separate least squares regressions of X and XV on Z and V , without taking into account that the model for X implies the model for XV , and without considering that the postulated models may be incompatible (e.g., even when the model for X includes a main effect of V , the model for XV may not allow for a main effect of V^2). Two-stage estimators (including TSLS estimators) are moreover generally inefficient when the true exposure relation is nonlinear in Z or C (e.g. because it includes an interaction between Z and components of C , or because it is of the logistic form), or when the outcome is dichotomous so that $\text{Cov}((X, Y)^T | Z, C)$ is not constant in Z, C . In particular, it may happen under certain data laws that the Standard TSLS estimator does not exist (more precisely, is not \sqrt{n} -consistent), even though other two-stage estimators with small variance exist. This is for instance the case when $E(X|Z, C) = Z - ZV$ for a scalar variate $V \in C$ which takes the values 0 and 1 with probability 1/2, independently of Z , and when furthermore $m(C; \psi)X = \psi X$ and $m_y(C; \beta) = \beta_0 + \beta_1 V$. In that case, the implied first stage model would ignore the interaction between Z and V and thus result in $E(X|Z, C) = 0$, thereby violating the necessary rank condition for TSLS estimators. In those cases, a Plug-In TSLS estimator based on a first stage model that includes main effects of Z , V and their interaction, is indicated.

4 Estimation without reliance on an exposure model

It follows from the proof of Proposition 2 that simultaneous fitting of the exposure and outcome model may be needed in order to obtain a semi-parametric efficient estimator of ψ^* in model $\mathcal{M} \cap \mathcal{A}_x \cap \mathcal{A}_y$. However, as this estimator may lack robustness against

misspecification of \mathcal{A}_x , we consider and recommend more robust procedures in this and subsequent sections. Before we address an estimation procedure that does not require an exposure model we note that certain Plug-in TSLS estimators of ψ^* in model \mathcal{M}^{const} enjoy robustness against misspecification of \mathcal{A}_x , in particular the Standard TSLS estimator (Robins, 2000; Wooldridge, 2002), even though they were not designed to this end. The basic, and maybe somewhat paradoxical, rationale is that for particular choices of \mathcal{A}_x it turns out that the estimating equations remain valid even if \mathcal{A}_x is misspecified while for other choices of \mathcal{A}_x the estimating equations then lose their validity.

We can see the robustness by noting that solving (12) will be equivalent to solving

$$0 = \sum_{i=1}^n e_y(Z_i, C_i) \{Y_i - m_y(C_i; \beta) - m(C_i; \psi)X_i\} \quad (13)$$

whenever the index function $e_x(Z, C)$ in (11) for fitting the exposure model \mathcal{A}_x includes the component

$$e_y(Z, C)m(C; \psi), \quad (14)$$

for all ψ (or some full-rank linear transformation of it). This is because the fitting procedure for the exposure model then ensures that

$$\sum_{i=1}^n e_y(Z_i, C_i)m(C_i; \psi)X_i = \sum_{i=1}^n e_y(Z_i, C_i)m(C_i; \psi)m_x(Z_i, C_i; \hat{\alpha}) \quad (15)$$

for all ψ . Estimating equation (13) no longer involves the exposure model. In particular, its unbiasedness is not dependent upon (correct) specification of an exposure model. For arbitrary functions $e_y(Z, C)$ of the dimension of (β, ψ) , the solution to equation (13) is thus a CAN estimator of (β^*, ψ^*) in model $\mathcal{M} \cap \mathcal{A}_y$, i.e., regardless of (correct) specification of model \mathcal{A}_x . In Section 4.1, we use the above result to show that certain two-stage estimators exhibit robustness against misspecification of the exposure model. In Section 4.2, we derive the semi-parametric efficient estimator under model $\mathcal{M} \cap \mathcal{A}_y$ (i.e., the optimal index function $e_y(Z, C)$ in (13)).

4.1 Robustness of two-stage estimators against misspecification of the exposure model

Condition (15) is met by certain two-stage estimators in linear models that are fitted using ordinary least squares, as detailed below.

Proposition 3 Robustness of Standard TSLS estimators against exposure model misspecification (see also Robins (2000) and Wooldridge (2002, Theorem 5.1))

The Standard TSLS estimator of ψ^ is CAN under model $\mathcal{M} \cap \mathcal{A}_y^{lin}$. This estimator is therefore still consistent when the implied exposure model \mathcal{A}_x^{lin} is not a correct model for $E(X|Z, C)$.*

While Proposition 3 is stated for Standard TSLS, it can be somewhat extended to other two-stage estimators as discussed in Appendix A of the Supplementary Materials. For instance, robustness can be attained by ensuring that the exposure model is linear and *minimally* includes the covariates in the outcome model. The above result does however, not extend to general two-stage estimators. It fails for plug-in estimators when the exposure effect depends on covariates, e.g. when $m(C; \psi) = \psi^T C$, or when the exposure model is nonlinear. For instance, consider fitting model $E(Y|Z, C) = \beta^* C + m_x(Z, C; \alpha^*) \psi^{*T} C$ by ordinary least squares. Because ordinary least squares uses index function $e_y(Z, C) = (C^T, m_x(Z, C; \alpha^*) C^T)^T$, we have that $e_y(Z, C) m(C; \psi)$ is then no longer contained in $e_x(Z, C)$ to ensure identity (15).

4.2 Local efficiency without exposure model

It is not necessary to rely on two stage estimators being ‘accidentally’ robust towards misspecification of \mathcal{A}_x (i.e. satisfying the conditions of Proposition 3); we can instead estimate ψ^* straightaway by solving an estimating equation of the form (13) yielding a

CAN estimator under $\mathcal{M} \cap \mathcal{A}_y$ as stated below. Without relying on \mathcal{A}_x , we may however lose efficiency compared to two-stage estimation when \mathcal{A}_x is correctly specified. Hence, in order to achieve greatest efficiency possible under $\mathcal{M} \cap \mathcal{A}_y$, the following Proposition 4 also addresses the optimal choice of the index function $e_y(Z, C)$ in (13).

Proposition 4 Semi-parametric efficient CAN estimation under model $\mathcal{M} \cap \mathcal{A}_y$

The IV estimator of the causal parameter ψ^ obtained by solving (13) for some conformable index function $e_y(Z, C)$ is CAN under $\mathcal{M} \cap \mathcal{A}_y$; it does not rely on an exposure model. Moreover, all CAN estimators of ψ^* in model $\mathcal{M} \cap \mathcal{A}_y$ are asymptotically equivalent to the solution of (13) for some conformable index function $e_y(Z, C)$. A semiparametric locally efficient estimator of (β^*, ψ^*) is obtained by choosing $e_y(Z, C)$ equal to*

$$e_{y,\text{opt}}(Z, C) = \begin{pmatrix} m_x(Z, C; \alpha^*) \partial m(C; \psi^*) / \partial \psi \\ \partial m_y(C; \beta^*) / \partial \beta \end{pmatrix} \text{Var}\{Y - m(C; \psi^*)X | Z, C\}^{-1};$$

the efficiency is local in the sense that it depends on specification of a working model \mathcal{A}_x , and is only attained when model $\mathcal{M} \cap \mathcal{A}_y \cap \mathcal{A}_x$ is correctly specified and the conditional variance $\text{Var}\{Y - m(C; \psi^)X | Z, C\}$ is consistently estimated (at faster than $n^{1/4}$ rate).*

To illustrate the above, let $m_x(Z, C; \alpha) = \text{expit}(\alpha_z^T Z + \alpha_c^T C)$, $m_y(C; \beta) = \beta^T C$ and $m(C; \psi) = \psi^T C$. Then under homoscedasticity (i.e. $\text{Var}\{Y - m(C; \psi^*)X | Z, C\}$ is constant) a semi-parametric efficient estimator of ψ^* in model $\mathcal{M} \cap \mathcal{A}_y$ is obtained by first estimating α^* using standard logistic regression, and next solving

$$0 = \sum_{i=1}^n \begin{pmatrix} \text{expit}(\hat{\alpha}_z^T Z_i + \hat{\alpha}_c^T C_i) C_i \\ C_i \end{pmatrix} (Y_i - \beta^T C_i - \psi^T C_i X_i);$$

by the linearity of this equation in ψ , a semi-parametric (locally) efficient estimator is thus obtainable in closed form.

Wooldridge (2002) also studies semi-parametric efficiency under model $\mathcal{M} \cap \mathcal{A}_y$, but his results relate to the subclass of estimators obtained by solving (13) for $e_y(Z, C)$ linear

in Z and C . He shows that Standard TSLS estimators are efficient within this class under model $\mathcal{M} \cap \mathcal{A}_y$. Likewise, efficiency results for GMM estimators (Wooldridge, 2002) relate to a subclass of all estimators under model $\mathcal{M} \cap \mathcal{A}_y$ (namely, the solutions to arbitrary linear combinations, with constant coefficients, of a given number of unbiased estimating equations). It follows from the above expression for the efficient score that potentially more efficient estimators than TSLS estimators can be obtained under model $\mathcal{M}^{const} \cap \mathcal{A}_y$ when the exposure is nonlinear in Z or C , or when there is heteroscedasticity.

In Appendix A of the Supplementary Materials we compare more generally the efficiency of the approach in Proposition 4 to that achieved using semi-parametric efficient joint estimation under a correct model for \mathcal{A}_x . We find that the loss in efficiency by not relying on correct specification of an exposure model is low when the confounding is weak or under \mathcal{A}_x^{lin} as implied by TSLS. Conversely, efficiency can potentially be gained when confounding is strong and e.g. \mathcal{A}_x is not linear.

5 Double-robust estimation

As the estimators of Proposition 4 do not rely on correct specification of an exposure model their validity is merely predicated upon correct specification of the structural model \mathcal{M} in the absence of covariates (i.e., when C is empty). When covariate adjustment is necessary, either because the IV assumptions are only satisfied conditional on covariates, or because of interest in effect heterogeneity, then they do rely on an outcome model, \mathcal{A}_y . Misspecification of that model may then sometimes result in biased effect estimates.

Robustness against misspecification of the outcome model \mathcal{A}_y is achieved by further restriction to a subclass of estimators for ψ^* obtained by solving (13) with $e_y(Z, C)$ an arbitrary vector function whose first p components, with p the dimension of ψ , equal

$$e(Z, C) - E\{e(Z, C)|C\}$$

for some arbitrary function $e(Z, C)$. Here, the conditional expectation $E\{e(Z, C)|C\}$ is calculated under a parametric model \mathcal{A}_z defined by

$$f(Z|C) = f(Z|C; \gamma^*),$$

where $f(Z|C; \gamma)$ is a known density function, smooth in γ , and γ^* is an unknown finite-dimensional parameter, which can be substituted by its maximum likelihood estimator $\hat{\gamma}$. For instance, when Z is binary, we may assume that $P(Z = 1|C) = \text{expit}(\gamma^{*T}C)$ and use standard logistic regression to estimate γ^* . Further, let $\hat{\beta}$ be a consistent estimator of β^* as obtained in the previous section. Then an estimator of ψ^* can be obtained by solving

$$0 = \sum_{i=1}^n [e(Z_i, C_i) - E\{e(Z_i, C_i)|C_i; \hat{\gamma}\}] \left\{ Y_i - m_y(C_i; \hat{\beta}) - m(C_i; \psi)X_i \right\}, \quad (16)$$

for some conformable vector function $e(Z, C)$.

Proposition 5 CAN estimation under model $\mathcal{M} \cap (\mathcal{A}_y \cup \mathcal{A}_z)$

The IV estimator of the causal parameter ψ^ obtained by solving (16) for some conformable index function $e(Z_i, C_i)$ is CAN under $\mathcal{M} \cap (\mathcal{A}_y \cup \mathcal{A}_z)$. Moreover, all CAN estimators of ψ^* in model $\mathcal{M} \cap (\mathcal{A}_y \cup \mathcal{A}_z)$ are asymptotically equivalent to the solution of (16) for some conformable index function $e(Z_i, C_i)$.*

Proof: see Robins (2000) and Okui et al. (2012). \square

Because the solution to (16) is a CAN estimator of ψ^* when either working model \mathcal{A}_z or \mathcal{A}_y holds, in addition to the linear IV model \mathcal{M} , it has been called double-robust (Robins and Rotnitzky, 2001; Okui et al., 2012). The resulting estimators, which are also known as G-estimators (Robins, 1994), are especially attractive in studies where the law of Z given C is known as this guarantees robustness against misspecification of \mathcal{A}_y . Such knowledge, leading to correct specification of \mathcal{A}_z , is for instance given in randomized experiments where Z denotes randomization, or in Mendelian randomization

studies where the genetic instrument is often known to be independent of covariates C , in which case $E\{e(Z, c)|C = c\}$ can be consistently estimated as $n^{-1} \sum_{i=1}^n e(Z_i, c)$. Note that typical two-stage estimators fail to exploit such knowledge of the law of Z given C .

5.1 Robustness of two-stage estimators against misspecification of the exposure and outcome model

The double-robustness property can be used to show that misspecification of the outcome model does not result in biased exposure effect estimates for the Standard TSLS estimator of ψ^* when the IV Z happens to be linear in the covariates of the outcome model (or, in particular, independent of C) in the sense that $E(Z|C) = \gamma^{*T}C$ (Robins, 2000; Okui et al., 2012); this is interesting as Standard TSLS appears not to make use of \mathcal{A}_z .

Proposition 6 Robustness of TSLS estimators against outcome model misspecification

The Standard TSLS estimator of ψ^ in model \mathcal{M} with $m(C; \psi)$ linear in C is CAN under model \mathcal{M} when $E(Z|C) = \gamma^{*T}C$.*

It follows from Proposition 6 that when Z and C are not independent, the robustness of the Standard TSLS estimator does not extend to general IVs, e.g. dichotomous IVs that obey a logistic regression model with main covariate effect C , nor to general two-stage estimators that involve nonlinear exposure models or effect heterogeneity (i.e. $m(C; \psi^*)$ depending on C). It further follows from the proof of Proposition 6 in the Supplemental Materials that the Plug-in TSLS estimator of ψ^* is CAN under model \mathcal{M}^{const} when the conditional mean $E\{m_x(Z, C; \alpha^*)|C\}$ is linear in C , and in the more general model \mathcal{M} with $m(C; \psi)$ linear in C when Z is independent of C , but not necessarily otherwise. Thus, when $E(Z|C)$ is linear in V and V^2 (with $C = (1, V)^T$), then the Plug-in TSLS estimator will only be robust against outcome model misspecification when the outcome model includes the term V^2 (regardless of whether it is associated with the outcome).

5.2 Efficiency of double-robust and TSLS estimators

When choosing $e(Z, C)$ in (16) one may want to consider the efficiency of the resulting estimator and ask whether it is worthwhile including covariates at all when there is the choice. To address this, we first recall how a semi-parametric (locally) efficient estimator of ψ^* under model $\mathcal{M} \cap \mathcal{A}_z$ is obtained. It follows from Robins (1994) (see also Okui et al., 2012) that $e(Z, C)$ in (16) should then be equal to

$$e_{\text{opt}}(Z, C) = \sigma^{-2}(Z, C) \frac{\partial m(C; \psi^*)}{\partial \psi} \left[E(X|Z, C) - \frac{E\{\sigma^{-2}(Z, C)E(X|Z, C)|C\}}{E\{\sigma^{-2}(Z, C)|C\}} \right] \quad (17)$$

with $\sigma^2(Z, C) \equiv \text{Var}\{Y - m(C; \psi^*)X|Z, C\}$. Since model $\mathcal{M} \cap (\mathcal{A}_z \cup \mathcal{A}_y)$ is less restrictive, this is also delivering the (locally) efficient estimator of ψ^* in model $\mathcal{M} \cap (\mathcal{A}_z \cup \mathcal{A}_y)$. For instance, assuming that $E(X|Z, C) = \alpha_1^{*T}C + \alpha_2^{*T}ZC$ for scalar Z and C , $m_y(C; \beta) = \beta^T C$ and $\sigma^2(Z, C) = \sigma^2$ for unknown parameters α_1^*, α_2^* and β^* , we have

$$e_{\text{opt}}(Z, C) = \sigma^{-2} \alpha_2^{*T} C \{Z - E(Z|C)\}. \quad (18)$$

A locally efficient G-estimator may now be obtained by substituting α_2^* by the ordinary least squares estimator in the above expression, setting σ^2 to 1 (as it is just a proportionality constant), and next solving (16) for the resulting choice of $e(Z, C) = e_{\text{opt}}(Z, C)$. These expressions suggest a way to optimally include covariates and, in a similar vein, to optimally combine multiple instruments (see e.g. Bowden and Vansteelandt, 2011).

In the special case where $Z \perp\!\!\!\perp C$ and under $\mathcal{M}^{\text{const}}$, one has the choice of whether to adjust for C at all. Consistent estimation can then also be achieved ignoring C . However, provided $\mathcal{M}^{\text{const}}$ and working models for $E(Y - \psi^*X|C)$, $E(X|Z, C)$ and $\sigma^2(Z, C)$ are correctly specified, the covariate-adjusted analysis will then be at least as efficient in large samples as the unadjusted analysis. While an efficiency gain is not generally guaranteed when these working models are misspecified, the following corollary demonstrates that it *can* (almost) always be guaranteed for the special case of Standard TSLS.

Corollary 7 Efficiency of covariate adjusted Standard TSLS estimators

*When $Z \perp\!\!\!\perp C$ and under \mathcal{M}^{const} , if $Y - \psi^*X$ is conditionally independent of Z given C , then covariate adjustment does not increase (and usually reduces) the asymptotic variance of the Standard TSLS estimator of ψ^* .*

Proof: this follows as a special case of Proposition 8; the proof is given in Appendix B of the Supplemental Materials. \square

The condition that $Y - \psi^*X$ be conditionally independent of Z given C in the above corollary can be regarded as a stronger version of the assumption of ‘no effect modification’ by Z common to IV models (Hernan and Robins, 2006; Clarke and Windmeier, 2010).

Fisher-Lapp and Goetghebuer (1999) also noticed that covariate adjustment is typically beneficial in a linear IV context; however, their results are specific to the case of partial compliance with full compliance in the control arm, where by design there is a corresponding interaction in the exposure model and where the IV model is specific to the treatment arm.

5.3 Simulation study

To better appreciate the above results, we show empirical results from a small simulation study with $n = 500$ in Figure 1 (a more extensive simulation study follows in Section 7). We generated mutually independent and standard normal covariates U and V , Z dichotomous with $P(Z = 1|U, V) = 0.27$, X dichotomous with $P(X = 1|Z, U, V) = \Phi(Z + U + V)$ and Y normal with mean $0.5X - U - 2V + V^2$. Assuming a linear IV model with $m(C; \psi) = \psi$, a logistic model for the IV with $P(Z = 1|C) = \text{expit}(\alpha_0 + \alpha_1 V)$ and $C = (1, V)^T$, we then evaluated the following estimators:

(TSLS): the Standard TSLS using a misspecified outcome models in that it excludes V^2 and, implicitly, a misspecified linear exposure model;

(TS): the plug-in two-stage estimator with $P(X = 1|Z, C) = \Phi(\alpha_z Z + \alpha_c^T C)$ (under which we obtained an average Nagelkerke pseudo- R^2 value of 0.24 and an F-value of 78) and $E(Y|Z, C) = \beta^T C + \psi\Phi(\hat{\alpha}_z Z + \hat{\alpha}_c^T C)$ which also wrongly excludes V^2 ;

(LE-y-c) and (LE-y-m): two locally efficient estimators under model $\mathcal{M} \cap \mathcal{A}_y$ as introduced in Section 4; the first with correct outcome model, the second with incorrect outcome model (i.e. $m_y(C; \beta) = \beta^T C$), but both using the correct exposure model;

(DR-cc, DR-cm and DR-mm): three double-robust semiparametric locally efficient estimators under model $\mathcal{M} \cap (\mathcal{A}_y \cup \mathcal{A}_z)$ (see Section 5); the first (DR-cc) uses correctly specified exposure and outcome model, the second (DR-cm) uses a wrong outcome model (see above), the third (DR-mm) a misspecified exposure model (i.e. $m_x(Z, C; \alpha) = \alpha_z Z + \alpha_c^T C$) and again the misspecified outcome model.

The results are displayed in Figure 1 (top). They reveal, as predicted by Propositions 3 and 5, that Standard TSLS (TSLS) is robust against misspecification of the exposure and outcome model when $Z \perp\!\!\!\perp C$, unlike the plug-in two-stage estimator (TS). As predicted in Section 4, the efficiency of Standard TSLS estimators can be improved by acknowledging that the exposure model is nonlinear (the relative efficiency of estimator LE-y-c versus TSLS is 0.38). However, the misspecified version (LE-y-m) is seriously biased. The double-robust estimators (DR-cc, DR-cm and DM-mm), are somewhat less efficient as they avoid reliance on a correct outcome model, (e.g. the relative efficiency of estimator LE-cc versus LE-y-c is 1.14). This is even more pronounced under model misspecification, because semi-parametric efficiency is only guaranteed when the exposure and outcome models are correctly specified. However, by their double robustness, these estimators remain unbiased under such misspecification as they correctly exploit that $Z \perp\!\!\!\perp C$.

Figure 1 (bottom) shows results from a setting where the IV depends nonlinearly on covariates, and where the degree of model misspecification is more pronounced. In particular, data were generated and analysed as before, but with $P(Z = 1|C, V) = \text{expit}(-1 + V/2)$ and $P(X = 1|Z, U, V) = \Phi(Z + U + V - ZV + V^2/2)$, corresponding to a Nagelkerke pseudo- R^2 value of 0.13 and an F-value of 38 under the working model $P(Z = 1|C) = \text{expit}(\alpha_0 + \alpha_1 V)$. The findings are similar to before, but the nonlinear dependence of the instrument on covariates impairs the robustness of the Standard TSLS estimator, as predicted by Proposition 5 (its bias is 0.55).

We also evaluated the behavior of the above estimators in the presence of effect modification. Data were generated as in the first simulation experiment above, but with X being normal with mean $2Z + V + U - ZV + 0.5V^2$ and Y normal with mean $0.5X + XV - U - 2V + V^2$. We evaluated the above estimators, but considered the Standard TSLS estimator with IVs Z and VZ , and with $C = (1, V, V^2)$ (TSLS-c) and with $C = (1, V)$ (TSLS-m) to evaluate the impact of correct model specification. In addition, we considered the two-stage estimator with $C = (1, V, V^2)$ (TS-c) and with $C = (1, V)$ (TS-m) in the outcome model, but using a misspecified linear exposure model with main effects of Z and V only. The results are analogous as before. Figure 2 (left) confirms the lack of robustness of two-stage estimators against misspecification of the exposure model, as opposed to Standard TSLS estimators (Figure 2, right). The best performance is seen for the locally efficient estimator under model $\mathcal{M} \cap \mathcal{A}_y$, but this estimator lacks robustness against misspecification of the outcome model. The Standard TSLS estimators and the double-robust estimators are much less efficient, especially under model misspecification. In the following sections, we will propose strategies to further improve performance.

Figure 1 and 2 about here.

6 Improved double-robust estimation

Consistency of the double-robust estimator of ψ^* demands correct specification of either the outcome model \mathcal{A}_y or the IV model \mathcal{A}_z ; local efficiency demands correct specification of both these models, and additionally of models for the exposure distribution and conditional outcome variance. In practice all these models are typically somewhat misspecified. In Section 6.1, we therefore propose a strategy to guarantee efficiency within a subclass of double-robust estimators as soon as the IV model \mathcal{A}_z is correctly specified. In Section 6.2, we propose strategies that aim to minimise locally the bias of the double-robust estimator when both the outcome model \mathcal{A}_y and the IV model \mathcal{A}_z are misspecified. Throughout these sections, results are confined to the main effect structural model \mathcal{M}^{const} .

6.1 Empirical efficiency maximisation

The semi-parametric efficient estimator of ψ^* , obtained by substituting the conditional expectations in (17) by estimates under parametric models, is not guaranteed to outperform simpler CAN estimators (e.g. obtained by solving (16) for $e(Z, C) = Z$ or by ignoring covariate information) under model misspecification, as we will see in the simulation study of Section 7. Okui et al. (2012) proposed regression double-robust estimators that have no larger asymptotic variance than a given double-robust estimator, even under model misspecification. In this Section, we generalise their results with the potential for bigger efficiency gains in return. We will realise this by building on and extending the ideas behind empirical efficiency maximisation, a procedure originally proposed by Rubin and van der Laan (2008) and Cao, Tsiatis and Davidian (2009) in the missing data literature. In this subsection, we assume that model \mathcal{A}_z is correctly specified.

Let $\hat{\psi}(\alpha, \beta)$ be the double-robust estimator of ψ^* obtained by solving estimating equation (16) for a user-specified parameterisation $e(Z, C; \alpha)$ of $e(Z, C)$, evaluated at the

given values α (and β indexing $m_y(C; \beta)$). This may, but need not be guided by the form of the efficient index function given in (17). For instance, for a scalar Z , one may postulate that $e(Z, C)$ is of the form $\alpha^T CZ$ for some α . When the law of Z given C is known, then the asymptotic variance of $\hat{\psi}(\alpha, \beta)$ under model $\mathcal{M} \cap \mathcal{A}_z$ equals

$$\frac{\text{Var}([e(Z, C; \alpha) - E\{e(Z, C; \alpha)|C\}]\{Y - m_y(C; \beta) - \psi^*X\})}{nE([e(Z, C; \alpha) - E\{e(Z, C; \alpha)|C\}]X)^2}. \quad (19)$$

Let $\tilde{\alpha}$ and $\tilde{\beta}$ be the values of α and β , respectively, that minimise the empirical analog of (19) with ψ^* substituted by a preliminary consistent estimator under model $\mathcal{M}^{const} \cap \mathcal{A}_z$, e.g. a G-estimator based on $e(Z, C) = Z$ and model \mathcal{A}_y^{lin} . The proposition below then shows that $\hat{\psi}(\tilde{\alpha}, \tilde{\beta})$ is a double-robust estimator which is at least as efficient as $\hat{\psi}(\alpha, \beta)$ for arbitrary α and β . Key properties that underlie the validity of the proposition are (a) that $\tilde{\beta}$ is CAN for β^* under model \mathcal{A}_y ; and (b) that $\hat{\psi}(\tilde{\alpha}, \tilde{\beta})$ and $\hat{\psi}(\tilde{\alpha}^*, \tilde{\beta}^*)$ have the same asymptotic variance under model $\mathcal{M}^{const} \cap \mathcal{A}_z$, with $\tilde{\alpha}^*$ and $\tilde{\beta}^*$ being the probability limits of $\tilde{\alpha}$ and $\tilde{\beta}$ (provided $\tilde{\alpha}$ and $\tilde{\beta}$ converge at faster than $n^{1/4}$ rate).

Proposition 8 Efficiency within a subclass of double-robust estimators

Let $\tilde{\alpha}$ and $\tilde{\beta}$ minimise the empirical version of (19). Then the estimator $\hat{\psi}(\tilde{\alpha}, \tilde{\beta})$ solving (16) is CAN under model $\mathcal{M}^{const} \cap (\mathcal{A}_y \cup \mathcal{A}_z)$.

Moreover, when the law of Z given C is known, then we have that for all α and β

$$\lim_{n \rightarrow \infty} \text{Var} \left[\sqrt{n} \left\{ \hat{\psi}(\tilde{\alpha}, \tilde{\beta}) - \psi^* \right\} \right] \leq \lim_{n \rightarrow \infty} \text{Var} \left[\sqrt{n} \left\{ \hat{\psi}(\alpha, \beta) - \psi^* \right\} \right].$$

Proof: see Appendix B. \square

In Appendix B we further discuss the case where the law of Z given C is known only up to a finite-dimensional parameter.

Consider for instance the choices $e(Z, C; \alpha) = \alpha^T CZ$ and $m_y(C; \beta) = \beta^T C$. Then by construction, $\hat{\psi}(\tilde{\alpha}, \tilde{\beta})$ is at least as efficient as the estimator obtained by solving (16) for the

simple choices $e(Z, C) = Z$ and $m_y(C) = 0$, i.e. the estimator which ignores covariates. Hence, when $Z \perp\!\!\!\perp C$, then the resulting approach will deliver a covariate adjustment strategy that is guaranteed to be at least as efficient as an unadjusted analysis. More generally, efficiency is - by construction - always attained within the subclass of estimators allowed by the models for $e(Z, C)$ and $m_y(C)$, but semi-parametric efficiency under model $\mathcal{M} \cap (\mathcal{A}_z \cup \mathcal{A}_y)$ is only attained when the efficient index function (17) happens to equal $e(Z, C; \alpha)$ for some α and when $E(Y - \psi^* X | C)$ equals $m_y(C; \beta)$ for some β .

Minimising the empirical analog of (19) can generally be done numerically, but in special cases also by suitably modified regression techniques. For example, we show in Appendix B of the Supplemental Materials that when $e(Z, C) = \alpha^T C Z$, then under certain assumptions minimising (19) w.r.t. α can be done by fitting the regression model $E(X | Z, C) = \alpha^T C \{Z - E(Z | C)\}$ using ordinary least squares. Minimising (19) w.r.t. β is possible by fitting the regression model $E(Y - \psi^* X | C) = \beta^T C$ using weighted least squares with weights $(\alpha^T C)^2 \{Z - E(Z | C)\}^2$. The above procedure needs some modification when the law of Z given C is unknown and the model \mathcal{A}_y is (possibly) misspecified.

The regression double-robust estimator of Okui et al. (2012) may be viewed as a special case of the above proposal. It fixes α at some given value (which may not minimise the asymptotic variance) and chooses $m_y(C; \beta) = \beta m_y(C)$ for some given $m_y(C)$.

6.2 Bias-reduced double-robust estimation

The efficiency results of Section 6.1 are especially attractive when model \mathcal{A}_z is known to hold, as is the case in certain study designs. In other cases, bias becomes, arguably, a more dominant concern. Although there seems little hope that one can avoid bias in the estimation of ψ^* when both working models \mathcal{A}_z and \mathcal{A}_y are misspecified, Vermeulen and Vansteelandt (2015) found that for quite a general class of double-robust estimators,

surprisingly, the nuisance parameters indexing \mathcal{A}_z and \mathcal{A}_y can be estimated so as to target bias reduction. Briefly, they note that the asymptotic bias (Stefanski and Boos, 2002) of an estimator for ψ^* , evaluated at fixed nuisance parameters β and γ , equals the expected value of its influence function $U(\psi^*, \beta, \gamma)$; for given α , this is here:

$$U(\psi, \beta, \gamma) = \frac{[e(Z, C; \alpha) - E\{e(Z, C; \alpha)|C; \gamma\}]\{Y - m_y(C; \beta) - \psi X\}}{E\{[e(Z, C; \alpha) - E\{e(Z, C; \alpha)|C; \gamma\}]X\}}.$$

Minimising the squared bias in the direction of β thus amounts to setting the gradient

$$2E\{U(\psi^*, \beta, \gamma)\}E\left\{\frac{\partial U}{\partial \beta}(\psi^*, \beta, \gamma)\right\}$$

to zero. Although the first component cannot generally be made zero without knowing aspects of the data-generating law, interestingly, the second component delivers an unbiased estimating function for γ (Vermeulen and Vansteelandt, 2015). This is so because, by the double-robustness, $U(\psi^*, \beta, \gamma)$ is mean zero for all β at γ^* when model \mathcal{A}_z holds. The second component can thus be made zero empirically, by using it as a basis for estimation, as illustrated in the next paragraph. Under local misspecification of one of the working models (as formally defined in Vansteelandt et al. (2012)), this procedure reduces the order of the asymptotic bias. Under gross misspecification, it prevents inflation of the asymptotic bias, although one cannot exclude that other nuisance parameter estimators happen to deliver less biased effect estimators under some data-generating mechanisms.

In the remainder of this section, we apply the bias-reduction procedure of Vermeulen and Vansteelandt (2015) to double-robust estimators in model $\mathcal{M}^{const} \cap (\mathcal{A}_z \cup \mathcal{A}_y)$. For illustration, suppose that the instrument Z is dichotomous with working model $P(Z = 1|C; \gamma) = \text{expit}(\gamma^T C)$, that $m(C; \beta) = \beta^T C$, and let the index function $e(Z, C; \alpha)$ be of the form $Ze(C; \alpha)$ for some $e(C; \alpha)$ (as is the case for the efficient score for ψ^* under model $\mathcal{M}^{const} \cap (\mathcal{A}_z \cup \mathcal{A}_y)$ when $E(X|Z, C)$ is linear in Z and $\text{Var}(Y|Z, C)$ does not depend on Z).

Taking the gradient of $U(\psi, \beta, \gamma)$ with respect to β then results in estimating equations

$$0 = \sum_{i=1}^n \frac{\partial U_i(\gamma, \beta)}{\partial \beta} = \sum_{i=1}^n e(C_i; \alpha) \{Z_i - P(Z_i = 1|C_i; \gamma)\} C_i, \quad (20)$$

which are unbiased for γ . Since γ and β are of the same dimension, γ can thus be estimated as the solution to this equation. Solving equation (20) ensures that

$$\sum_{i=1}^n e(C_i; \alpha) \{Z_i - P(Z_i = 1|C_i; \gamma)\} m_y(C_i; \beta) = 0$$

so that the estimating equation for ψ reduces to

$$\begin{aligned} 0 &= \sum_{i=1}^n e(C_i; \alpha) \{Z_i - P(Z_i = 1|C_i; \gamma)\} \{Y_i - m_y(C_i; \beta) - m(C_i; \psi^*)X_i\} \\ &= \sum_{i=1}^n e(C_i; \alpha) \{Z_i - P(Z_i = 1|C_i; \gamma)\} \{Y_i - m(C_i; \psi^*)X_i\} \end{aligned}$$

which no longer involves β . Bias-reduced estimation of γ then overcomes the need to estimate β , and thereby prevents that the choice of estimator of β amplifies bias. Solving (20) may not be straightforward for certain data sets. We therefore adapt the proposal of Vermeulen and Vansteelandt (2015) by extending the logistic regression model for Z to

$$P(Z = 1|C; \gamma) = \text{expit} \{ \gamma^T C + \theta^T C e(C; \alpha) \}.$$

This model contains the original working model \mathcal{A}_z (corresponding to $\theta = 0$). Moreover, fitting this model using the default maximum likelihood procedure has the effect of making the identity (20) hold, as the latter corresponds with the score for the coefficient of $e(C; \alpha)C$. The resulting procedure will be referred to as BR- γ .

When using the procedure BR- γ , we continue to estimate α indexing $e(C; \alpha)$ as explained in Section 6.1. Although now, we no longer assume that model \mathcal{A}_z is correctly specified, estimating α in this manner still has the effect of minimising the asymptotic variance of the double-robust estimator across all values of α . This is because the procedure BR- γ sets the gradient of the influence function w.r.t. β equal to zero, so that there

is no need to account for the estimation of β in the calculation of the asymptotic variance (Vermeulen and Vansteelandt, 2015). Because BR- γ moreover employs maximum likelihood estimation under a particular extended model for the IV, conservative standard errors can be obtained as 1 over root n times the empirical standard deviation of $U(\gamma, \beta)$, ignoring the estimation of γ (Rotnitzky, Li and Li, 2010). Alternatively, robust sandwich standard errors can be calculated, or the bootstrap can be used.

We also considered a related approach whereby we estimated γ using maximum likelihood and β by setting the gradient of the influence function $U(\gamma, \beta)$ with respect to γ to zero. This results in the following unbiased estimating equations for β :

$$0 = \sum_{i=1}^n \frac{\partial U_i(\psi^*, \gamma, \beta)}{\partial \gamma} = \sum_{i=1}^n e(C_i; \alpha) \{Y_i - m_y(C_i; \beta) - \psi^* X_i\} \Gamma_i \quad (21)$$

for

$$\begin{aligned} \Gamma_i = & \{Z_i - P(Z_i = 1|C_i; \gamma)\} E[e(C_i; \alpha)P(Z_i = 1|C_i; \gamma)P(Z_i = 0|C_i; \gamma)C_i X_i] \\ & - P(Z_i = 1|C_i; \gamma)P(Z_i = 0|C_i; \gamma)C_i E[e(C_i; \alpha) \{Z_i - P(Z_i = 1|C_i; \gamma)\} X_i] \end{aligned}$$

It can be verified that the effect of the factor Γ_i is to eliminate ψ^* from the estimating equation so that knowledge of the truth ψ^* is not needed for estimating β . This approach is designed to locally minimise the bias of the double-robust estimator in the direction of γ , at the maximum likelihood estimate $\hat{\gamma}$. To solve (21), we jointly fit an extended linear model for the outcome $Y - \psi^* X$ with covariates C and $e(C; \alpha)P(Z = 1|C; \gamma)P(Z = 0|C; \gamma)C$ using ordinary least squares (where, again, the choice of ψ^* does not affect results), and the (double-robust) estimating equation for ψ . This has the effect of making the identity (21) hold. The resulting procedure will be referred to as BR- β . By setting the gradient of the influence function $U(\gamma, \beta)$ with respect to γ equal to zero, it need not adjust for the estimation of γ (Vermeulen and Vansteelandt, 2015). However, the uncertainty in the estimate of β must be acknowledged using sandwich standard errors or the bootstrap.

7 Simulation study

We conducted a simulation experiment with $n = 500$ independent measurements on mutually independent and standard normal covariates U and V , Z dichotomous with $P(Z = 1|V) = \text{expit}(-1 + V/2 + \lambda_z V^2/3)$, X normal with mean $Z + U + V - ZV + \lambda_x V^2$ and Y normal with mean $X - U - V + \lambda_y V^2$. Assuming a linear IV model with $m(C; \psi) = \psi$, we then evaluated the following estimators:

1. TSLS: the Standard TSLS estimator using $(Z, VZ)^T$ as IV vector, based on a linear model for the exposure, involving main effects of V, Z and their interaction, and a linear model for the outcome involving main effects of V and the fitted value from the first stage regression. Including the VZ interaction in the first stage model ensures a fairer comparison with the subsequent estimators so that for all estimators misspecification of the exposure model is only due to omitting V^2 .
2. Loc Eff: the locally efficient double-robust estimator (assuming homoscedasticity) based on a logistic model for Z with a main effect of V , a linear model for X with a main effect of Z and V and their interaction, and a linear outcome model (i.e., $m_y(C) = \beta^T C$ with $C = (1, V)^T$), all fitted using maximum likelihood.
3. Emp Eff: the locally efficient double-robust estimator using the same fitted model for Z as before, but using working models $e(Z, C) = \alpha^T C Z$ and $m_y(C) = \beta^T C$ fitted using empirical efficiency maximization (ignoring estimation of the model for Z , which is suboptimal when the outcome model is misspecified).
4. BR- β , BR- γ : the double-robust estimator with α^* estimated using empirical efficiency maximization, but with either the outcome model or the IV model fitted using bias-reduced estimation.

To obtain the estimators Loc Eff and Emp Eff, the TSLS estimator was used as a starting value; the obtained estimate was then updated a single time. In the calculation of BR- β , BR- γ was used as a starting value as it was easy to obtain and generally performing well.

Table 1 shows the simulation results based on 1000 simulations. When all working models are correctly specified (i.e. $\lambda_z = \lambda_x = \lambda_y = 0$), then all estimators have nearly identical performance to Standard TSLS. This is theoretically expected, because they are based on correctly specified working models in the calculation of the efficient score and are therefore asymptotically equivalent. When only the outcome model is misspecified (i.e. $\lambda_z = \lambda_x = 0, \lambda_y \neq 0$), then the TSLS estimator is biased (as the instrument distribution does not satisfy Proposition 6) with larger standard deviation than the double-robust estimators, which were all unbiased. Bias-reduced estimation of the outcome model (BR- β) resulted in major efficiency gains in this case. When only the exposure model was misspecified (i.e. $\lambda_z = \lambda_y = 0, \lambda_x \neq 0$), then as theoretically predicted, the TSLS estimator and the double-robust estimators continue to be unbiased, but the performance of the locally efficient double-robust estimator was sometimes very poor because its efficiency is only attained at a correctly specified model for the exposure. In this case, drastic improvements were obtained via empirical efficiency maximization, because this strategy guarantees efficiency within a subclass of estimators, regardless of correct specification of an exposure model. The efficiency of the resulting double-robust estimator was sometimes better, sometimes worse than that of TSLS. When only the IV model was misspecified (i.e. $\lambda_x = \lambda_y = 0, \lambda_z \neq 0$), then all estimators were unbiased because of the double-robustness of the estimators and the fact that the TSLS estimator does not rely on correct specification of an IV model; all estimators had nearly identical performance in this case. When both the exposure and outcome model are misspecified (i.e. $\lambda_z = 0, \lambda_x \neq 0, \lambda_y \neq 0$), then again TSLS is biased, unlike the double-robust estimators. The locally efficient es-

estimator behaved poorly in this case and is greatly outperformed by empirical efficiency maximisation, which again performs best in combination with bias-reduced estimation of the outcome model. When all models were misspecified, then also the double-robust estimators were subject to bias. However, bias-reduced estimation of either the outcome model or the IV model resulted in bias reductions and efficiency gains. This is not surprising for BR- γ because the extended IV model happened to contain the truth: indeed, the inclusion of the covariate $e(C; \alpha)^T C$ in the instrument model was tantamount to the inclusion of V^2 . For BR- β , where the extended outcome model did not contain the truth, this confirms that the proposed procedure reduces bias under model misspecification.

To further evaluate the bias-reduced estimation strategy, we additionally ran simulations under extreme misspecifications, such that both extended outcome and IV models did not contain the truth. In particular, we generated $n = 500$ independent measurements on mutually independent and standard normal covariates U and V , Z dichotomous with $P(Z = 1|V) = 1 - \exp\{-\exp(-1 + V/2 - V^2/2 + \lambda_z V^2/8)\}$, X normal with mean $Z + U + V - ZV + 2V^2 + 2ZV^2 + 2\lambda_x V^3$ and Y normal with mean $X - U - V - 2V^2 + 2\lambda_y V^3$.

The working models were the same as before. The results are visualised in Figure 3 for all combinations of λ_x, λ_y and λ_z in $\{-1, 1\}$, and confirm the previous findings. The locally efficient double-robust estimator had very poor performance and, while empirical efficiency maximization resulted in major efficiency gains, it was still much worse than TSLS estimation. For instance, in the setting of Figure 3 (top, left), the locally efficient double-robust estimator had bias and standard deviation of -32.7 and 450, as opposed to -0.61 and 4.1 with empirical efficiency maximization, and -0.54 and 3.1 with TSLS. In combination with bias-reduced estimation, most of the bias disappeared and variance was often greatly reduced (see Figure 3).

Table 1 and Figure 3 about here.

8 Illustration

We illustrate the proposed methodology on a sample of 3010 working men aged between 24 and 34 who were part of the 1976 wave of the US National Longitudinal Survey of Young Men (Card, 1995). In particular, we will estimate the effect of years of education on the log of hourly wages in 1976 (Y). Following Card (1995), we use as an IV an indicator if the individual lived close to a college that offered 4 year courses in 1966 (Z). All reported analyses are adjusted for covariates (C) years of labour market experience and its square, marital status, an indicator if the individual is black, as well as various measures of geographical location in 1966 and 1976. Twelve years of education was most common (33%) in this study and was therefore used as a reference class by defining X to be the difference between the years of education and 12.

The log of hourly wages is reasonably normally distributed with mean 6.3 (SD 0.44), and is on average 0.075 (95% CI 0.068 to 0.082) higher per extra year of education, after linear regression adjustment for years of labour market experience, marital status, race and geographical location in 1966 and 1976. The partial correlation between education and the IV is 0.066. Below, we will report the results from IV analysis with 95% percentile-based confidence intervals based on the nonparametric bootstrap with 1000 resamples.

Standard TSLS analysis yields an education effect of 0.13 (SE 0.067, 95% CI 0.029 to 0.28) on the average log of the hourly wage, corresponding with a one-year increase in education. Because the instrument is dichotomous and strongly associated with covariates, its expectation is likely nonlinear in the covariates. The Standard TSLS estimator is therefore sensitive to correct specification of the role of covariates in the outcome model. We thus evaluate the double-robust estimators based on a logistic regression model for the IV. The locally efficient G-estimator equals 0.10 (SE 0.044, 95% CI 0.025 to 0.18). Like

the double-robust estimator based on empirical efficiency maximization (0.088, SE 0.045, 95% CI 0.0063 to 0.18), it is much more efficient than the Standard TSLS estimator. Further, more minor efficiency gains are obtained through the proposed bias reduction strategies. In particular, we find that $\text{BR-}\gamma$ equals 0.092 (SE 0.041, 95% CI 0.010 to 0.18) and $\text{BR-}\beta$ equals 0.095 (SE 0.043, 95% CI 0.0063 to 0.19).

9 Discussion

In this article, we have argued that Standard TSLS estimation, unlike many variations of the two-stage approach to estimation with an IV, is often robust against misspecification of the working models for the exposure and outcome. However, this robustness may come at the expense of a loss of precision, which can be considerable when, for instance, the exposure mean is nonlinear in the instrument and/or covariates, e.g. when the exposure is binary, multinomial or count data. Moreover, the suggested robustness of the Standard TSLS estimators is limited to specific data-generating mechanisms: robustness against misspecification of the outcome model is for instance lost in Standard TSLS estimators when the IV is nonlinear in covariates. We also demonstrated that another strength of Standard TSLS, not generally shared by other two-stage estimators, is that including covariates will asymptotically not reduce, and typically improve, efficiency when instrument and covariates are known to be independent and in the absence of effect modification.

In contrast, locally efficient double-robust IV estimators confer robustness against model misspecification in a wider class of data generating mechanisms. For instance, an attractive alternative, when instruments and covariates are known to be independent, is the estimator obtained by empirical efficiency maximisation: it is guaranteed consistent and efficient relative to a subclass of all CAN estimators. In other situations one should arguably worry more about bias than efficiency. We have shown that major improve-

ments can be achieved by combining empirical efficiency maximisation with bias-reduced double-robust estimation. The resulting estimators have a very stable performance with considerable robustness against misspecification of all models for the instrument, exposure and outcome; their standard errors can be computed relatively easily using sandwich estimators. We are hopeful that by extending these results to double-robust estimators in nonlinear IV models (Robins, 1994; Vansteelandt et al., 2010), we will be able to improve the performance of IV estimators in these more complex settings where difficulties of estimation are common (Vansteelandt et al., 2011; Burgess et al., 2014).

There are some limitations to our work. Our results are asymptotic and do not take into account the problem of ‘weak instrument / small sample’ bias (Bound et al., 1995). This may in practice exacerbate the problem of bias due to model misspecification. There are a number of variations on two-stage estimators that are designed to address this problem, such as e.g. limited information maximum likelihood (Anderson, 2004), but these will not generally exhibit comparable robustness towards model misspecification. It would be an important area for future research to tackle both sources of bias simultaneously. Related to this, although the results on empirical efficiency maximisation appear to suggest that it is beneficial to adjust for all available covariates C when $Z \perp\!\!\!\perp C$, the performance of the resulting estimators may be affected in the presence of high-dimensional covariates. Whether and how to best select covariates in such cases, as well as in settings where it is not known whether $Z \perp\!\!\!\perp C$, constitutes an important area for future research.

References

- Anderson, T.W. (2005). Origins of the limited information maximum likelihood and two-stage least squares estimators. *Journal of Econometrics*, **127**, 1-16.
- Angrist, J.D., Imbens, G.W., and Rubin, D.B. (1996). Identification of causal effects

- using instrumental variables. *Journal of the American Statistical Association*, **91**, 444-455.
- Bound, J., Jaeger, D.A. and Baker, R.M. (1995). Problems with instrumental variables estimation when the correlation between the instruments and the endogenous variable is weak. *Journal of the American Statistical Association*, **90**, 443-50.
- Bowden, J. and Vansteelandt, S. (2011). Mendelian randomisation analysis of case-control data using Structural Mean Models. *Statistics in Medicine*, **30**, 678-694.
- Bowden, R.J. and Turkington, D.A. (1985). *Instrumental Variables*. Cambridge University Press.
- Burgess, S., Granel, R., Palmer, T.M., Sterne, J.A.C. and Didelez, V. (2014). Lack of identification in semi-parametric instrumental variable models with binary outcomes. *American Journal of Epidemiology*,
- Card, D. (1995). *Using geographic variation in college proximity to estimate the return to schooling*. In L. Christophides, E. Grant and R. Swidinsky (eds.), *Aspects of Labour Market Behaviour: Essays in Honour of John Vanderkamp*.
- Cao, W.H., Tsiatis, A.A. and Davidian, M. (2009). Improving efficiency and robustness of the double-robust estimator for a population mean with incomplete data. *Biometrika*, **96**, 723-734.
- Chamberlain, G. (1987). Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics*, **34**, 305-334.
- Clarke, P.S. and Windmeijer, F. (2010). Identification of causal effects on binary outcomes using structural mean models. *Biostatistics*, **11**, 756-770.
- Didelez, V. and Sheehan, N.A. (2007). Mendelian randomization as an instrumental variable approach to causal inference. *Statistical Methods in Medical Research*, **16**, 309-330.

- Didelez, V., Meng, S. and Sheehan, N.A. (2010). Assumptions of IV methods for observational epidemiology. *Statistical Science*, **25**, 22-40.
- Fischer-Lapp, K., and Goetghebeur, E. (1999). Practical properties of some structural mean analyses of the effect of compliance in randomized trials. *Controlled Clinical Trials*, **20**, 531-546.
- Greenland, S. (2000). An introduction to instrumental variables for epidemiologists. *International Journal of Epidemiology*, 29 (4), 722-729.
- Hernán, M.A. and Robins, J.M. (2006). Instruments for causal inference - An epidemiologist's dream? *Epidemiology*, **17**, 360-372.
- Kleibergen, F. and Zivot, E. (2003). Bayesian and classical approaches to instrumental variable regression. *Journal of Econometrics*. **114**, 29-72.
- Mullahy, J. (1997). Instrumental variable estimation of count data models: Application to models of cigarette smoking behaviour. *Review of Economics and Statistics*, 79, 586-593.
- Okui, R., Small, D.S., Tan, Z.Q. and Robins, J.M. (2012). Doubly robust instrumental variable regression. *Statistica Sinica*, **22**, 173-205.
- Pearl, J. (2009). *Causality: Models, Reasoning and Inference*. Cambridge University Press, New York, NY, USA, 2nd edition.
- Robins, J.M. (1994). Correcting for non-nompliance in randomized trials using structural nested mean models. *Communications in Statistics: Theory and Methods* **23**, 2379-2412.
- Robins JM. (2000). Robust estimation in sequentially ignorable missing data and causal inference models. *Proceedings of the American Statistical Association*, Section on Bayesian Statistical Science 1999, pp. 6-10.

- Robins, J.M., and Rotnitzky, A. (2001). Comment on “Inference for semiparametric models: Some questions and an answer,” by P.J. Bickel and J. Kwon. *Statistica Sinica*, **11**, 920-936.
- Rotnitzky, A., Li, L.L., and Li, X.C. (2010). A note on overadjustment in inverse probability weighted estimation. *Biometrika*, **97**, 997-1001.
- Rubin, D.B. and van der Laan, M.J. (2008). Empirical efficiency maximization: improved locally efficient covariate adjustment in randomized experiments and survival analysis. *International Journal of Biostatistics*, **4**, 1-40.
- Stefanski, L.A. and Boos, D.D. (2002). The calculus of M-estimation. *The American Statistician*, **56**, 29-38.
- Tchetgen Tchetgen, E.J., Walter, S., Vansteelandt, S., Martinussen, T. and Glymour, M. (2015). Instrumental variable estimation in a survival context. *Epidemiology*, **26**, 402-410.
- Vansteelandt, S., Bowden, J., Babanezhad, M. and Goetghebeur, E. (2011). On instrumental variables estimation of causal odds ratios. *Statistical Science*, **26**, 403-422.
- Vansteelandt, S., Bekaert, M. and Claeskens, G. (2012). On model selection and model misspecification in causal inference. *Statistical Methods in Medical Research*, **21**, 7-30.
- Vermeulen, K. and Vansteelandt, S. (2015). Bias-reduced doubly robust estimation. *Journal of the American Statistical Association*, in press.
- Wooldridge, J. (2002). *Econometric Analysis of Cross Section and Panel Data*, MIT Press, Cambridge, MA.




Figure 1: Simulation results for binary exposure. Top: instrument independent of covariates; Bottom: instrument nonlinear in covariates.




Figure 2: Simulation results under effect modification. Boxplots of the main exposure effect estimators (top) and the exposure-covariate interaction estimators (bottom). Some of the outlying estimates are not visualised.

Table 1: Empirical bias and standard deviation of the two-stage estimator (TS), the locally efficient double-robust estimator (Loc Eff), the double-robust estimator based on empirical efficiency maximization (EEM) and these same estimators that employ bias-reduced nuisance parameter estimators (BR). The superscript number between brackets refers to the number of severely outlying estimates that were eliminated in the calculation of bias and empirical standard deviation.

	λ_x	λ_y	λ_z	TS	Loc Eff	EEM	BR- β	BR- γ
Bias	0	0	0	0.0033	0.0043	0.0044	0.0041	0.0042
	0	1	0	-0.55	-0.0092	-0.035	0.0024	-0.017
	0	-1	0	0.56	0.018	0.044	0.0059	0.026
	1	0	0	0.000073	0.013	0.0058	0.0037	0.0046
	-1	0	0	0.0013	0.0074	0.0043	0.0048	0.0044
	0	0	1	-0.00057	-0.00033	$9.0 \cdot 10^{-5}$	$-9.1 \cdot 10^{-5}$	0.00029
	0	0	-1	0.0053	0.0055	0.0095	0.0050	0.0045
	1	1	0	0.15	$0.11^{(2)}$	-0.040	0.0039	-0.019
	-1	1	0	-0.41	0.012	-0.021	0.0016	-0.013
	1	-1	0	-0.15	$-0.095^{(4)}$	0.051	0.0038	0.028
	-1	-1	0	0.41	0.0030	0.030	0.0079	0.022
	1	1	1	0.34	0.36	0.11	0.021	-0.00028
	-1	1	1	-0.35	$-14^{(2)}$	-0.40	0.024	0.00073
	1	-1	1	-0.34	-0.36	-0.11	-0.023	-0.00059
	-1	-1	1	0.35	$15^{(2)}$	0.86	-0.023	0.0015
	1	1	-1	-0.94	$0.59^{(1)}$	-0.48	0.019	0.0057
	-1	1	-1	-0.36	-0.084	-0.085	0.018	0.0048
	1	-1	-1	0.94	$-0.20^{(2)}$	0.50	-0.0081	0.0039
	-1	-1	-1	0.37	0.10	0.10	-0.0086	0.0036
SD	0	0	0	0.11	0.11	0.11	0.11	0.11
	0	1	0	0.24	0.18	0.17	0.12	0.17
	0	-1	0	0.28	0.19	0.19	0.12	0.18
	1	0	0	0.18	$0.82^{(4)}$	0.12	0.12	0.12
	-1	0	0	0.068	0.12	0.11	0.11	0.11
	0	0	1	0.094	0.097	0.11	0.097	0.097
	0	0	-1	0.13	0.13	0.14	0.13	0.13
	1	1	0	0.46	$1.9^{(2)}$	0.19	0.12	0.17
	-1	1	0	0.11	0.23	0.17	0.12	0.17
	1	-1	0	0.48	$1.2^{(4)}$	0.20	0.12	0.18
	-1	-1	0	0.14	0.22	0.17	0.12	0.17
	1	1	1	0.15	0.12	0.16	0.10	0.14
	-1	1	1	0.15	$120^{(2)}$	8.30	0.11	0.14
	1	-1	1	0.14	0.10	0.16	0.099	0.13
	-1	-1	1	0.16	$140^{(2)}$	12	0.10	0.13
	1	1	-1	1.3	$11^{(1)}$	0.85	0.13	0.17
	-1	1	-1	0.098	$40^{(1)}$	0.18	0.14	0.17
	1	-1	-1	1.5	$5.6^{(2)}$	0.98	0.13	0.17
	-1	-1	-1	0.12	0.17	0.18	0.13	0.16

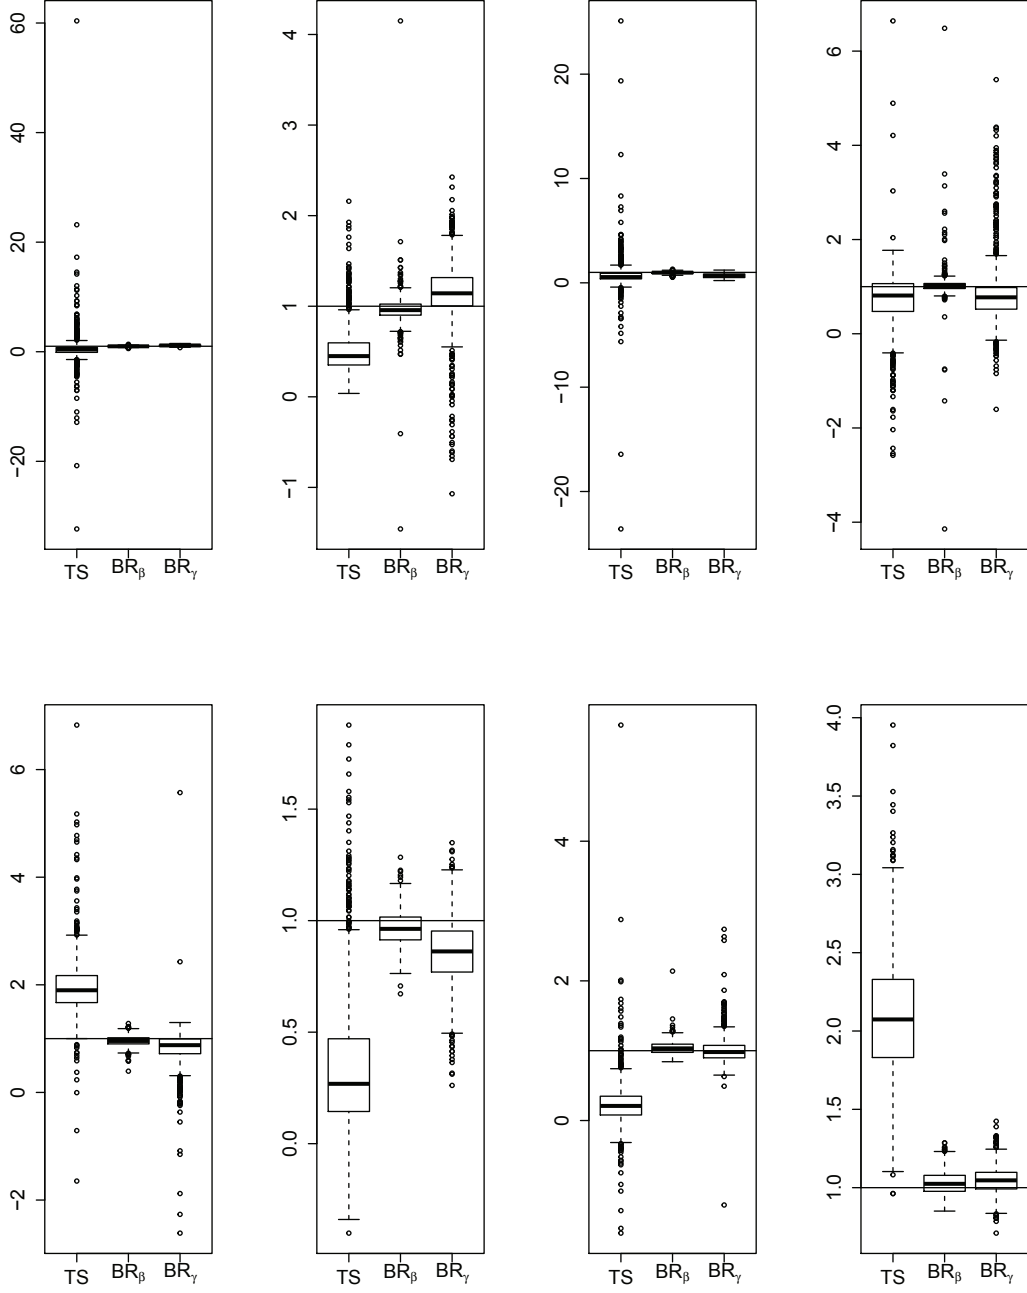


Figure 3: Boxplots of the two-stage estimator (TS), the double-robust estimator based on empirical efficiency maximization and bias-reduced nuisance parameter estimators (BR) under extreme model misspecification.